



Configuring DHCP Failover

DHCP failover is a protocol designed to allow a backup DHCP server to take over for a main server if the main server is taken off the network for any reason. You can use DHCP failover to configure two DHCP servers to operate as a redundant pair.

Failover Scenarios


There are three basic failover scenarios:

- Simple failover—One server acting as main and its partner acting as backup.
- Back office failover—Two mains having the same backup server.
- Symmetrical failover—Two servers acting as main and backup for each other.

Simple Failover

Simple failover involves a main server and a single backup server pair (see [Figure 16-1](#)). In the example, main server A has three scopes that must be configured identically on backup server B.

Figure 16-1 Simple Failover Example



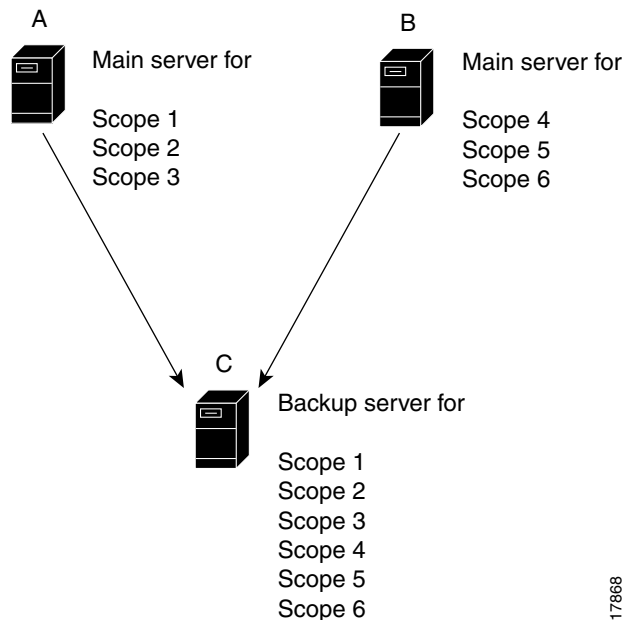
The advantages of simple failover over the other scenarios are:

- It is the easiest to manage as the network changes—It is fully supported by the Web UI so that changes to the main server configuration are automatically propagated to the backup server.
- Provides the greatest performance benefits.
- You only need to set the failover properties on the server level and not worry about the scopes.

Back Office Failover

Back office failover involves two (or more) main servers that share the same backup server (see [Figure 16-2](#)). In the example, main servers A and B have different scopes, and backup server C must include all these scopes. This scenario is appropriate for scopes on the same LAN segment, which require the same main and backup servers, but with the sets of scopes on different LAN segments.

Figure 16-2 Back Office Failover Example




An advantage of back office failover over the other scenarios is that it reduces the number of servers managed. However, simple failover is still recommended, because in back office failover:

- The backup server must be sized to handle the sum of the configurations.
- Changes to any of the main servers must be duplicated on the backup server.
- The increased complexity of the configuration management can substantially reduce the actual availability of the configuration.

Symmetrical Failover

Symmetrical failover involves servers that act as backups for each other (see [Figure 16-3](#)). This scenario is extremely tricky in that there can be no variance in scope attribute values between the servers, or the relationship will not work properly.

Figure 16-3 Symmetrical Failover Example

The disadvantage to symmetrical failover over the other scenarios is that, while reducing the number of servers, there is little to no performance benefit. A backup server operates at about 40% of the main server to keep its lease database synchronized. If the servers back each other up, a portion of their processing capacity goes to this task, with less capacity available to servicing clients. Moreover, because each scope must be configured individually, symmetrical failover is more prone to configuration errors. Because of these significant disadvantages, simple failover is the recommended method.

Failover Checklist

Use this checklist to prepare for an effective failover configuration:

- Duplicate the scope, policy, DHCP option, and address configurations on the partner servers—The Network Registrar Web UI provides a way to automate this process.
- Ensure that both partners are configured with a wide enough range of addresses so that the backup server can provide leases while the main server is down for a reasonable amount of time.
- If you change any of the following configurations on the main server, also change them on the backup server:
 - Scopes, including ensuring identical scope-selection tags
 - Policies
 - IP addresses
 - Reservations
 - Clients
 - Client-classes
 - Dynamic DNS updates
 - Dynamic BOOTP
 - Virtual private networks (VPNs)
 - DHCP extensions
- If you use LDAP, direct the partner servers to the same LDAP server.

- If you use BOOTP relay (IP helpers), configure all BOOTP relay agents to point to both partners. Network Registrar does not automatically detect this. You can only detect BOOTP configuration errors by performing live tests in which you periodically take the main server out of service to verify that the backup server is available to DHCP clients.

Configuring Failover for DHCP Servers and Scopes

You can use the Network Registrar Web UI or CLI to configure DHCP failover server pairs. The types of configuration options supported by managing failover server pairs are:

- Policy properties and DHCP options, including vendor-specific options
- DHCP server properties
- Scope properties and ranges
- Reservations
- Clients and client-classes
- Scope selection tags
- Extensions

To add a failover pair, you must set the failover attributes on the DHCP server or scope level:

Configuring Failover on DHCP Servers

One of two failover configuration methods is to set up failover on the server level. You basically enable the local DHCP server for failover, then set which servers you want as the main and backup for the failover pair. You would generally use this method for simple failover configurations (see the [“Simple Failover”](#) section on page 16-1).

Step 1 Enable the local DHCP server for failover:

- In the local cluster Web UI—On the Primary Navigation bar, click **DHCP**, then on the Secondary Navigation bar, click **DHCP Server** to open the Manage DHCP Server page. Click the Local DHCP Server link to open the Edit DHCP Server page. On this page, enable the following attribute:

Failover Setting—Click the **on** radio button.

- In the CLI—Use the **dhcp enable failover** command:

```
nrcmd> dhcp enable failover
```


Step 2 Set the main and backup servers (you must specify both) for the failover pair:

- In the local cluster Web UI—On the Edit DHCP Server page, set the following attributes:
 - Main Server—Enter the IP address of the main DHCP server in the failover pair.
 - Backup Server—Enter the IP address of the backup DHCP server in the failover pair.
- In the CLI—Use the **dhcp set failover-main-server** and **dhcp set failover-backup-server** commands to set the IP addresses of the main and backup servers. For example:

```
nrcmd> dhcp set failover-main-server=192.168.50.2
nrcmd> dhcp set failover-backup-server=192.168.60.2
```



Note Cisco recommends that you specify the IP address of the servers rather than their domain names, so that name resolution is not required for the failover pair to communicate.

- Step 3** Configure the other failover attributes as needed, such as the failover backup percentage and maximum client lead time (MCLT). The default values for these attributes are generally acceptable for most configurations. (See the “[Setting Advanced Failover Attributes](#)” section on page 16-14.)
- Step 4** Save the settings—In the local cluster Web UI, click **Modify Server**.
- Step 5** Reload the DHCP server:
- In the local cluster Web UI—On the DHCP Server page, click the Reload icon (.
 - In the CLI—Use the **dhcp reload** command:

```
nrcmd> dhcp reload
```

Configuring Failover for Scopes

The second of the two failover configuration methods is to set up failover for each applicable scope. This is known as scope-based failover and has similar settings as server-based failover. This method is generally used for back office or symmetrical failover, which are not recommended as configurations.

- Step 1** Create or edit a DHCP scope to have a wide enough range of IP addresses so that the backup server can provide leases while the main server is down for a reasonable amount of time. You can specify this range for the scope or, in the Web UI, for a scope template.
- In the local Web UI:
 - To add a new scope to enable failover—On the Secondary Navigation bar, click **Scopes** to open the List/Add DHCP Scopes page.
 - To edit an existing scope to enable failover—On the Secondary Navigation bar, click **Scopes** to open the List/Add DHCP Scopes page. Click the name of the existing scope to open the Edit DHCP Scope page.
 - To create a scope template and enable failover for the template for any scopes you create based on it—On the Secondary Navigation bar, click **Scope Templates** to open the List DHCP Scope Templates page. Click **Add Scope Template** to open the Add DHCP Scope Template page.
 - In the CLI—Use the **scope name create address mask** command to create the scope. You cannot create a scope template in the CLI.
- Step 2** There are two settings to enable failover for the scope (or scope template in the Web UI):
- *scope-enabled*—Choose this if you want to enable failover only for the scope and use the main and backup servers from the scope. You have to specify the main and backup servers for this setting.
 - *use-server-settings*—Choose this only if the server is also enabled for failover (see the “[Configuring Failover on DHCP Servers](#)” section on page 16-4) and you want to adopt the server settings for the scope. You do not need to specify the main and backup servers for this setting.

For further details on the scope failover settings, see the “[Scope Failover Attribute States](#)” section on page 16-14.

Apply these settings:

- In the local cluster Web UI—On the appropriate page, choose the *scope-enabled* or *use-server-settings* value from the Failover Setting drop-down list.
- In the CLI—Use the **scope name set failover=value** command, where *value* can be *scope-enabled* or *use-server-settings*. For example:

```
nrcmd> scope examplescope set failover=scope-enabled
```

Step 3 If appropriate, set the main and backup servers (you must specify both) for the failover pair:

- In the local cluster Web UI—Set the following attributes for the scope or scope template:
 - Main Server—Enter the IP address of the main DHCP server in the failover pair.
 - Backup Server—Enter the IP address of the backup DHCP server in the failover pair.
- In the CLI—Use the **scope name set failover-main-server** and **scope name set failover-backup-server** commands to set the IP addresses of the main and backup servers. For example:

```
nrcmd> scope examplescope set failover-main-server=192.168.50.2
nrcmd> scope examplescope set failover-backup-server=192.168.60.2
```



Note Cisco recommends that you specify the IP address of the servers rather than their domain names, so that name resolution is not required for the failover pair to communicate.

Step 4 Configure the other failover attributes as needed, such as the failover backup percentage and maximum client lead time (MCLT). The default values for these attributes are generally acceptable for most configurations. (See the “[Setting Advanced Failover Attributes](#)” section on page 16-14.)

Step 5 Save the settings—In the local cluster Web UI, click **Modify Scope**.

Step 6 Reload the DHCP server:

- In the local cluster Web UI—On the DHCP Server page, click the Reload icon (🔄).
- In the CLI—Use the **dhcp reload** command:

```
nrcmd> dhcp reload
```

Creating a Main and Backup Server Configuration

Once you enable a DHCP server or scope for failover, you can create a failover pair to include the server. A failover pair consists of a main and backup DHCP server. You can create failover pairs in the local and regional cluster Web UIs, and you can synchronize the main and backup servers. This functionality is not available in the CLI, which requires an explicit export to the backup server.

Creating and Synchronizing a Failover Pair on the Local Cluster

In the local cluster Web UI, when you enable failover on the local DHCP server (or scopes related to that server) and specify the main and backup servers, this automatically creates a failover pair. You must edit this failover pair to initiate connectivity between the servers. You then synchronize the failover pair so that the scopes, policies, clients, extensions and other DHCP properties match between the servers.

- Step 1** Create a failover pair in the local cluster Web UI—On the Primary Navigation bar, click **DHCP**, then on the Secondary Navigation bar, click **Failover** to open the List DHCP Failover Pairs page. (Note that the main and backup DHCP servers you specified when enabling failover (see the “[Configuring Failover for DHCP Servers and Scopes](#)” section on page 16-4) are now named as a hyphenation of the IP addresses of these server, and that the name is a link on the page.)
- Step 2** Click the failover pair name to open the Edit DHCP Failover Pair page.
- Step 3** On this page, you can see the main and backup server addresses, and whether the server pair is set up to use the attribute values from the server (true) or the scope (false) by default. Enter values in the following fields so that you can build the connectivity between the two servers:
- *remote-username*—Username to access the backup server. Required.
 - *remote-password*—Password to the backup server. Required.
 - *remote-scp-port*—CCM SCP port number to communicate with the target failover server. Check the target system for this port number, which is set during Network Registrar installation. On Windows systems, the installation sets the CNR_CCM_PORT registry key. On Solaris and Linux systems, the installation sets the CNR_CCM_PORT variable in the *install-dir/conf/cnr.conf* file. The default is **1234**. Required.
- Step 4** Click **Modify Failover**.
- Step 5** Click the Run icon (⊕) in the Synchronize column next to the pair name to open the Run Synchronize Failover Pair page. (If you click the Report icon (⊗), this opens the Report Synchronize Failover Pair page, which has the same contents.)
- Step 6** Choose the synchronization operation, depending on the degree to which you want the main server’s property values to replace those of the backup server. There are three basic operations:
- **Update**—This is the default and least radical operation. It is appropriate for update synchronizations in that it has the least effect on the unique properties of the backup server.
 - **Complete**—This operation is appropriate for all initial synchronizations. It is more complete than an update operation, while still preserving many of the backup server’s unique properties, such as are required for back office failover configurations.
 - **Exact**—This operation is appropriate for initial simple and symmetrical failover configurations, and is not appropriate for back office configurations. It makes the two servers as much as possible mirror images of each other, although it retains unique DHCP server, LDAP event services, and extension points on the backup server. (For initial failover configurations, use the Exact or Complete operation.)

Each operation performs a different mix of functions on the failover properties, as described in [Table 16-1](#). There are four functions, with examples based on these property name-value pairs:

On the main server:	On the backup server:
Name1=A	Name2=B
Name2=C	Name3=D





- **no change**—Makes no change to the list of properties or their values on the backup server. For the example, the result would be Name2=B, Name3=D.
- **ensure**—Ensures that a copy of the main server property exists on the backup server, but does not replace its value. For the example, the result would be Name1=A, Name2=B, Name3=D.
- **replace**—Replaces the value of a property that the two servers have in common with that of the main server. For the example, the result would be Name1=A, Name2=C, Name3=D.
- **exact**—Puts an exact copy of the main server’s list of properties and values on the backup server and removes the unique ones. For the example, the result would be Name1=A, Name2=C.

Table 16-1 Synchronization Functions Based on Update, Complete, or Exact Operations

Data Description	Update	Complete	Exact
DHCP Server (server level failover pair):	replace	replace	replace
Client-Class Properties			
Failover Properties			
Failover Tuning Properties			
Dynamic DNS Security Properties			
(See Table 16-2 for a list of the DHCP attributes affected by failover synchronization)			
All other Properties	no change	replace	replace
LDAP Event Service	no change	replace	replace
Policy:			
Option-list Property	ensure	replace	exact
All other Properties	replace	replace	exact
Client	replace	replace	exact
Client-Class	replace	replace	exact
Scopes (related to failover pair)	exact	exact	exact
VPN	replace	replace	exact
Key	replace	replace	exact
Extensions	ensure	replace	exact
Note You must manually copy over the extension files.			
Extension Point	no change	replace	replace
Option Information:	ensure	exact	exact
Custom options list			
Vendor options list			
Option-Data-types list			

Table 16-2 DHCP Attributes Affected by Failover Synchronization



Affected Failover Attributes	
<i>append-user-class-id-to-selection-tag</i>	<i>failover-main-server</i>
<i>client-class client-class-lookup-id</i>	<i>failover-poll-interval</i>
<i>defer-lease-extensions</i>	<i>failover-recover</i>
<i>dynamic-dns-fwd-key</i>	<i>failover-safe-period</i>
<i>dynamic-dns-rev-key</i>	<i>failover-use-safe-period</i>
<i>dynamic-dns-tsig</i>	<i>force-dns-updates</i>
<i>failover</i>	<i>map-user-class-id</i>
<i>failover-backup-percentage</i>	<i>skip-client-lookup</i>
<i>failover-backup-server</i>	<i>upgrade-unavailable-timeout</i>
<i>failover-dynamic-bootp-backup-percentage</i>	<i>use-ldap-client-data</i>
<i>failover-lease-period-factor</i>	<i>validate-client-name-as-mac</i>

- Step 7** Click **Run** on the Run Synchronize Failover page, or **Report** on the Report Synchronize Failover page:
- If you click **Run** and if the connection was accepted, the resulting View DHCP Failover Pair Sync Report page shows what change entries the synchronization added.
 - If you click **Report**, the resulting View DHCP Failover Pair Sync Report page shows what change entries the synchronization will apply if you run the synchronization. A **Run Update**, **Run Complete**, or **Run Exact** button indicates what kind of synchronization you want to perform. Click the applicable button to run the synchronization.
- Step 8** On the List DHCP Failover Pairs page, click the View icon () in the Manage Servers column to open the Manage DHCP Failover Servers page.
- Step 9** Click the Reload icon () next to the backup server to reload the backup server.
- Step 10** Try to get a lease.
- Step 11** On the Manage DHCP Failover Servers page, look at the health of the servers (they should show as ). Also, click the Logs icon () to view the log entries on the Log for Server page, and ensure that the servers are in NORMAL failover mode. The log file should contain an item similar to the following:
- ```
06/19/2003 9:41:19 name/dhcp/1 Info Configuration 0 04092 Failover is enabled
server-wide. Main server name: '192.168.0.1', backup server name: '192.168.0.110', mclt =
3600, backup-percentage = 10, dynamic-bootp-backup-percentage = 0, lease-period-factor =
150, use-safe-period: disabled, safe-period = 0.
```

## Creating and Synchronizing Failover Pairs on the Regional Cluster

In the regional cluster Web UI, you can create failover pairs and synchronize them with the servers on the local clusters. In this way, the local servers can start giving out leases in a failover environment.

- Step 1** Create and synchronize with the clusters that include the main and backup DHCP servers—On the Primary Navigation bar, click **Clusters**, then on the Secondary Navigation bar, click **Cluster List**. (See the “[Setting Up Server Clusters](#)” section on page 5-1.) This creates the subnets you can choose from which to draw the leases. You can also create the scope templates that you can push to the local clusters and use for the failover configuration.
- Step 2** Create a failover pair in the regional cluster Web UI—On the Primary Navigation bar, click **DHCP Configuration**, then on the Secondary Navigation bar, click **Failover** to open the List Failover Pairs page.
- Step 3** Click **Add Failover Pair** to open the Add Failover Pair page (see [Figure 4-22 on page 4-32](#)).
- Step 4** On this page:
- a. Give the failover pair a distinguishing name.
  - b. Choose the cluster for the main DHCP server from the drop-down list.
  - c. Choose the cluster for the backup server.
  - d. If you want to use a scope template to define the scopes for the failover pair, choose a scope template from the drop-down list.
  - e. Move the desired subnets over to the Selected field.
- Step 5** Make adjustments to the additional attributes, if needed. For a description, see the “[Setting Advanced Failover Attributes](#)” section on page 16-14.




- Step 6** Click **Add Failover Pair**. You return to the List DHCP Failover Pairs page. Because the connectivity between the servers is already set up through the cluster configuration, you do not need to specify them.
- Step 7** Click the Run icon () in the Synchronize column next to the name of the pair you just created to open the Run Synchronize Failover Pair page. If you click the Report icon () in the Synchronize column, this opens the Report Synchronize Failover Pair page, which has the same content as the Run Synchronize Failover Pair page.
- Step 8** Choose the synchronization operation, depending on the degree to which you want the main server's property values to replace those of the backup server. There are three basic operations:
- **Update**—Default and least radical operation. It is appropriate for update synchronizations in that it has the least effect on the unique properties of the backup server.
  - **Complete**—Appropriate for all initial synchronizations. It is more complete than an update operation, while still preserving many of the backup server's unique properties, such as are required for back office failover configurations.
  - **Exact**—Appropriate for simple and symmetrical failover configurations, and not for back office configurations. It makes the two servers as much as possible mirror images of each other, although it retains unique DHCP server, LDAP event services, and extension points on the backup server.






**Note** For initial failover configurations, use the **Exact** or **Update** operation.

Each operation performs a different mix of functions on the failover properties, as described in [Table 16-1 on page 16-8](#). There are four functions, with examples based on these property name-value pairs:

|                     |                       |
|---------------------|-----------------------|
| On the main server: | On the backup server: |
| Name1=A             | Name2=B               |
| Name2=C             | Name3=D               |


- **no change**—Makes no change to the list of properties or their values on the backup server. For the example, the result would be Name2=B, Name3=D.
  - **ensure**—Ensures that a copy of the main server property exists on the backup server, but does not replace its value. For the example, the result would be Name1=A, Name2=B, Name3=D.
  - **replace**—Replaces the value of a property that the two servers have in common with that of the main server. For the example, the result would be Name1=A, Name2=C, Name3=D.
  - **exact**—Puts an exact copy of the main server's list of properties and values on the backup server and removes the unique ones. For the example, the result would be Name1=A, Name2=C.
- Step 9** Click **Run** on the Run Synchronize Failover page, or **Report** on the Report Synchronize Failover page:
- If you click **Run** and if the connection was accepted, the resulting View DHCP Failover Pair Sync Report page shows what change entries the synchronization added.
  - If you click **Report**, the resulting View DHCP Failover Pair Sync Report page shows what change entries the synchronization will apply if you run the synchronization. A **Run Update**, **Run Complete**, or **Run Exact** button indicates what kind of synchronization you want to perform. Click the applicable button to run the synchronization.
- Step 10** On the List DHCP Failover Pairs page, click the Go Local icon () next to the main DHCP server for the failover pair to open a connection to the local cluster Web UI.
- Step 11** Go to the Manage Servers page on the local cluster and click the Reload icon () to reload the main server.
- Step 12** Click the Go Regional icon () to return to the regional cluster Web UI.

- Step 13** On the regional cluster List DHCP Failover Pairs page, click the Go Local icon () next to the backup DHCP server for the failover pair to open a connection to the backup server's local cluster Web UI, then reload the backup server.
- Step 14** Try to get a lease.
- Step 15** On the Manage DHCP Failover Servers page of the backup server's local cluster, look at the health of the servers (they should show as ). Also, click the Logs icon () to view the log entries on the Log for Server page, and ensure that the servers are in NORMAL failover mode. The log file should contain an item similar to the following:

```
06/19/2003 9:41:19 name/dhcp/1 Info Configuration 0 04092 Failover is enabled
server-wide. Main server name: '192.168.0.1', backup server name: '192.168.0.110', mclt =
3600, backup-percentage = 10, dynamic-bootp-backup-percentage = 0, lease-period-factor =
150, use-safe-period: disabled, safe-period = 0.
```

## Confirming Failover

Follow these steps to ensure that failover is enabled.

- Step 1** Ping from one server to the other to verify TCP/IP connectivity. Make sure that routers are configured to forward clients to both servers.
- Step 2** Check that the server is in NORMAL mode by clicking the Related Servers icon () on the Manage DHCP Server or List DHCP Failover Pairs page in the Web UI (see the [“Listing Related Servers for DHCP Servers”](#) section on page 5-3), or use the `dhcp getRelatedServers` command in the CLI.
- Step 3** After startup, have a client attempt to get a lease.
- Step 4** Set the log settings on the main server to include at least *failover-detail*.
- Step 5** Confirm that the `name_dhcp_1_log` log file (in *install-path/logs*) on the main server contains DHCPBNDACK or DHCPBNDUPD messages from each server.
- Step 6** Confirm that the `name_dhcp_1_log` log file on the backup server contains messages that the backup server is dropping requests because failover is in NORMAL state.
- Step 7** Repeat [Step 2](#).

## State Transitions During Integration

During normal operation, the failover partners transition between states. They stay in their current state until all the actions for the state transition are completed and, if communication fails, until the conditions for the next state are fulfilled. [Table 16-3](#) describes what happens when servers enter various states and how they initially integrate and later re-integrate with each other under certain conditions.

**Table 16-3 Failover State Transitions and Integration Processes**

| Integration                                                                  | Results                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |
|------------------------------------------------------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Into NORMAL state, the first time the backup server contacts the main server | <ol style="list-style-type: none"> <li>1. The newly configured backup server contacts the main server, which starts in PARTNER-DOWN state.</li> <li>2. Because the backup server is a new partner, it goes into RECOVER state and sends a Binding Request message to the main server.</li> <li>3. The main server replies with Binding Update messages that include the leases in its lease state database.</li> <li>4. After the backup server acknowledges these messages, the main server responds with a Binding Complete message.</li> <li>5. The backup server goes into RECOVER-DONE state.</li> <li>6. Both servers go into NORMAL state.</li> <li>7. The backup server sends Pool Request messages.</li> <li>8. The main server responds with the leases to allocate to the backup server based on the <i>failover-backup-percentage</i> configured.</li> </ol>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |
| After COMMUNICATIONS-INTERRUPTED state                                       | <ol style="list-style-type: none"> <li>1. When a server comes back up and connects with a partner in this state, the returning server moves into the same state and then immediately into NORMAL state.</li> <li>2. The partner also moves into NORMAL state.</li> </ol>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |
| After PARTNER-DOWN state                                                     | <p>When a server comes back up and connects with a partner in this state, the server compares the time it went down with the time the partner went into this state.</p> <ul style="list-style-type: none"> <li>• If the server finds that it went down and the partner subsequently went into this state: <ol style="list-style-type: none"> <li>a. The returning server moves into RECOVER state and sends an Update Request message to the partner.</li> <li>b. The partner returns all the binding data it was unable to send earlier and follows up with an Update Done message.</li> <li>c. The returning server moves into RECOVER-DONE state.</li> <li>d. Both servers move into NORMAL state.</li> </ol> </li> <li>• If the returning server finds that it was still operating when the partner went into PARTNER-DOWN state: <ol style="list-style-type: none"> <li>a. The server goes into POTENTIAL-CONFLICT state, which also causes the partner to go into this state.</li> <li>b. The main server sends an update request to the backup server.</li> <li>c. The backup server responds with all unacknowledged updates to the main server and finishes off with an Update Done message.</li> <li>d. The main server moves into NORMAL state.</li> <li>e. The backup server sends the main server an Update Request message requesting all unacknowledged updates.</li> </ol> </li> </ul> |

**Table 16-3 Failover State Transitions and Integration Processes (continued)**

| Integration                                     | Results                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
|-------------------------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
|                                                 | <ul style="list-style-type: none"> <li>f. The main server sends these updates and finishes off with an Update Done message.</li> <li>g. The backup server goes into NORMAL state.</li> </ul>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
| After the server loses its lease state database | <p>A returning server usually retains its lease state database. However, it can also lose it because of a catastrophic failure or intentional removal.</p> <ol style="list-style-type: none"> <li>1. When a server with a missing lease database returns with a partner that is in PARTNER-DOWN or COMMUNICATIONS-INTERRUPTED state, the server determines whether the partner ever communicated with it. If not, it assumes to have lost its database, moves into RECOVER state, and sends an Update Request All message to its partner.</li> <li>2. The partner responds with binding data about every lease in its database and follows up with an Update Done message.</li> <li>3. The returning server waits the maximum client lead time (MCLT) period, typically one hour, and moves into RECOVER-DONE state. For details on the MCLT, see the <a href="#">“Setting the Maximum Client Lead Time and Lease Period Factor”</a> section on page 16-17.</li> <li>4. Both servers then move into NORMAL state.</li> </ol> |
| After a lease state database backup restoration | <p>When a returning server has its lease state database restored from backup, and if it reconnects with its partner without additional data, it only requests lease binding data that it has not yet seen. This data may be different from what it expects.</p> <ol style="list-style-type: none"> <li>1. In this case, you must configure the returning server with the <i>failover-recover</i> attribute set to the time the backup occurred.</li> <li>2. The server moves into RECOVER state and requests all its partner’s data. The server waits the MCLT period, typically one hour, from when the backup occurred and goes into RECOVER-DONE state. For details on the MCLT, see the <a href="#">“Setting the Maximum Client Lead Time and Lease Period Factor”</a> section on page 16-17.</li> <li>3. Once the server returns to NORMAL state, you must unset its <i>failover-recover</i> attribute, or set it to zero.</li> </ol> <pre>nrcmd&gt; dhcp set failover-recover=0</pre>                                  |

Table 16-3 Failover State Transitions and Integration Processes (continued)

| Integration                                        | Results                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |
|----------------------------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| After the operational server had failover disabled | <p>If the operating server had failover enabled, disabled, and subsequently re-enabled, you must use special considerations when bringing a newly configured backup server into play. The backup server must have no lease state data and must have the <i>failover-recover</i> attribute set to the current time minus the MCLT interval, typically one hour. For details on the MCLT, see the “<a href="#">Setting the Maximum Client Lead Time and Lease Period Factor</a>” section on page 16-17.</p> <ol style="list-style-type: none"> <li>1. The backup server then knows to request all the lease state data from the main server. Unlike what is described in “After the server loses its lease state database” section of this table, the backup server cannot request this data automatically because it has no record of having ever communicated with the main server.</li> <li>2. After reconnecting, the backup server goes into RECOVER state, requests all the main server’s lease data, and goes into RECOVER-DONE state.</li> <li>3. Both servers go into NORMAL state. At this point, you must unset the backup server’s <i>failover-recover</i> attribute, or set it to zero.</li> </ol> <pre>nrcmd&gt; dhcp set failover-recover=0</pre> |

## Scope Failover Attribute States

The scope *failover* attribute has three possible states:

- **scope-enabled**—Indicates that this scope, and all scopes that are secondary to it or to which it is a secondary on this LAN segment, are enabled for failover. Scope parameters (not server parameters) should determine the main and backup servers.

If more than one scope on the same LAN segment is scope-enabled for failover, then the main and backup servers must be identical for each. An error occurs if one scope on a LAN segment is scope-enabled and another is scope-disabled, unless the other scope has failover enabled server-wide.

- **scope-disabled**—Disables a scope and all other scopes associated with it on a LAN segment from participating in failover. It only has meaning if failover is defined server-wide.
- **use-server-settings**—Indicates that this scope should use the settings for main and backup servers unless another scope associated with it on a LAN segment is either explicitly scope-enabled or scope-disabled. If one scope on a LAN segment is scope-enabled or scope-disabled, it overrides any scope for which *use-server-settings* is set on that LAN segment. Note that if you set the scope attribute *failover-back-percentage* explicitly, Network Registrar uses it, even if you set the *use-server-settings* attribute.

## Setting Advanced Failover Attributes

The advanced failover properties that can be important to set are the following:

- Backup percentage

- Maximum client lead time (MCLT) and lease period factor
- Safe period
- Request and response packet buffers
- Polling attributes
- Network discovery

## Setting Backup Percentages

To keep failover partners operating despite a network partition (when both servers can communicate with clients, but not with each other), allocate more addresses than for a single server. Configure the main server to allocate a percentage of the currently available addresses in each scope to the backup server. This makes these addresses unavailable to the main server. The backup server uses these addresses when it cannot talk to the main server and cannot tell if it is down.



### Note

If a Network Registrar failover server receives an update from a Network Registrar DHCP server running prior to Network Registrar 6.0, the unavailable leases do not have a timeout value. In this case, the Network Registrar 6.1 server uses the *unavailable-timeout* value configured in the scope policy or **system\_default\_policy** policy as the timeout for the unavailable lease. When the lease times out, the policy causes the lease to transition to available in both failover partners.

You can set the percentage of currently available addresses by setting the *failover-backup-percentage* attribute on the server or scope. Note that setting the backup percentage on the server level sets the value for all scopes not set with that attribute. However, if set at the scope level, the backup percentage overrides the one at the server level.

The backup percentage should be set large enough to allow the backup server to continue serving new clients in the event that the main server fails. The backup percentage is calculated based on the number of available addresses. The default backup percentage is 10%. However, this number can safely be set to a larger value, if extended outages are expected, because the main servers will periodically reclaim addresses (once per hour) if in the course of normal leasing activity, the main server's available address pool drops below its predefined percentage. For example, with the default 10% backup percentage, the main server will reclaim addresses if its address pool falls below 90%.

The percentage depends on the new client arrival rate and the network operator's reaction time. The backup server needs enough addresses from each scope to satisfy all new clients requests arriving during the time it does not know if the main server is down. Even during PARTNER-DOWN state, the backup server waits for the maximum client lead time (MCLT) and lease time to expire before re-allocating leases. See the [“Setting the Maximum Client Lead Time and Lease Period Factor”](#) section on page 16-17. When these times expire, the backup server offers:

- Leases from its private pool.
- Leases from the main server's pool.
- Expired leases to new clients.

During the day, an operator likely responds within two hours to COMMUNICATIONS-INTERRUPTED state to determine if the main server is working. The backup server then needs enough addresses to support a reasonable upper bound on the number of new clients that could arrive during those two hours.

During off-hours, the arrival rate of previously unknown clients is likely to be less. The operator can usually respond within 12 hours to the same situation. The backup server then needs enough addresses to support a reasonable upper bound on the number of clients that could arrive during those 12 hours.

The number of addresses over which the backup server requires sole control is the greater of the two numbers. You would express this number as a percentage of the currently available (unassigned) addresses in each scope. If you use client-classes, remember that some clients can only use some sets of scopes and not others.

**Note**

During failover, clients can sometimes obtain leases whose expiration times are shorter than the amount configured. This is a normal part of keeping the server partners synchronized. Typically this happens only for the first lease period, or during COMMUNICATIONS-INTERRUPTED state.

For all servers or scopes for which you enable failover, you must set the *failover-backup-percentage* attribute. This is the number of currently available (unreserved) leases that the backup server can use for allocations to new DHCP clients when the main server is down. You can use the default, which is 10 percent, or specify another value.

For scopes for which you enable dynamic BOOTP, use the *failover-dynamic-bootp-backup-percentage* attribute rather than the *failover-backup-percentage* attribute. The *failover-dynamic-bootp-backup-percentage* is the percentage of available addresses that the main server should send to the backup server for use with BOOTP clients.

The *failover-dynamic-bootp-backup-percentage* is distinct from the *failover-backup-percentage* attribute, because if you enable BOOTP on a scope, a server, even in PARTNER-DOWN state, never grants leases on addresses that are available to the other server. Network Registrar does not grant leases because the partner might give them out using dynamic BOOTP, and you can never safely assume that they are available again.

**Note**

You must define the dynamic BOOTP backup percentage on the main server. If you define it on the backup server, Network Registrar ignores it (to enable duplicating configuration through scripts). If you do not define it, Network Registrar uses the default *failover-backup-percentage*.

To properly support dynamic BOOTP while using the failover protocol, do this on every LAN segment in which you want BOOTP support:

- Create one scope for dynamic BOOTP.
- Enable BOOTP and dynamic BOOTP.
- Disable DHCP for that scope.

## Setting Backup Allocation Boundaries

You can be more specific as to which addresses to allocate to the backup server by using the *failover-backup-allocation-boundary* attribute on the scope. The IP address set as this value is the upper boundary of addresses from which to allocate addresses to a backup server. Only addressees below this boundary are allocated to the backup. If there are none available below this boundary, then the addresses above it, if any, are allocated to the backup. The actual allocation works down from this address, while the normal allocation for DHCP clients works up from the lowest address in the scope.

If you set *failover-backup-allocation-boundary*, you must also enable the *allocate-first-available* attribute on the scope. If *failover-backup-allocation-boundary* is unset or set to zero, then the boundary used is halfway between the first and last addresses in the scope ranges. If there are no available addresses below this boundary, then the first available address is used.

## Setting the Maximum Client Lead Time and Lease Period Factor

You can set two properties for failover that control certain adjustments to the lease period, the maximum client lead time (MCLT) and the lease period factor. These adjustments are essential for failover.

- **MCLT**—Controls the maximum allowed time beyond the expiration of a lease offered a client that the partner server knows the expiration to be. The default MCLT is one hour, which is optimized for most configurations. As defined by the failover protocol, the lease period given a client can never be more than the MCLT added to the most recently received potential expiration time from the failover partner, or the current time, whichever is later. That is why you sometimes see the initial lease period as only an hour, or an hour longer than expected for renewals. This hour is the MCLT, a form of lease insurance. The actual lease time is recalculated when the main server comes back.

The MCLT is necessary because of failover's use of *lazy updates*. Using lazy updates, the server can issue or renew leases to clients before updating its partner, which it can then do in batches of updates. If the server goes down and cannot communicate the lease information to its partner, the partner may try to re-offer the lease to another client based on what it last knew the expiration to be. The MCLT guarantees that there is an added window of opportunity for the client to renew. The way that a lease offer and renewal works with the MCLT is:

- a. The client sends a DHCPDISCOVER to the server, requesting a desired lease period (say, three days). The server responds with a DHCPPOFFER with an initial lease period of only the MCLT (one hour by default). The client then requests the MCLT lease period and the server acknowledges it.
- b. The server sends its partner a bind update containing the lease expiration for the client as the current time plus the MCLT. The update also includes the potential expiration time as the current time plus the client's desired period plus the MCLT (three days plus an hour). The partner acknowledges the potential expiration, thereby guaranteeing the transaction.
- c. When the client sends a renewal request halfway through its lease (in one-half hour), the server acknowledges with the client's desired lease period (three days). The server then updates its partner with the lease expiration as the current time plus the desired lease period (three days), and the potential expiration as the current time plus the desired period and another half of this period ( $3 + 1.5 = 4.5$  days). The partner acknowledges this potential expiration of 4.5 days. In this way, the main server tries to have its partner always lead the client in its understanding of the client's lease period so that it can always offer it to the client.

There is no one correct value for the MCLT. There is an explicit trade-off between various factors in choosing one. Most people use the default of one hour effectively and it works well in almost all environments. Here are some of the trade-offs between a short and long MCLT:

- **Short MCLT**—A short MCLT value means that after entering PARTNER-DOWN state, a server only has to wait a short time before it can start allocating its partner's IP addresses to DHCP clients. Furthermore, it only has to wait a short time after a lease expires before it can re-allocate that address to another DHCP client. However, the down side is that the initial lease interval that is offered to every new DHCP client will be short, which causes increased traffic, because those clients need to send their first renewal in a half of a short MCLT time. Also, the lease extensions that a server in COMMUNICATIONS-INTERRUPTED state can give is the MCLT only after the server has been in that state for around the desired client lease period. If a server stays in that state for that long, then the leases it hands out will be short and that will increase the load on that server, possibly causing difficulty.
- **Long MCLT**—A long MCLT value means that the initial lease period will be longer and the time that a server in COMMUNICATIONS-INTERRUPTED state can extend leases (after it being in that state for around the desired client lease period) will be longer. However, a server entering PARTNER-DOWN state must wait the longer MCLT before being able to allocate its partner's

addresses to new DHCP clients. This may mean that additional addresses are required to cover this time period. Also, the server in PARTNER-DOWN state must wait the longer MCLT from every lease expiration before it can re-allocate an address to a different DHCP client.

- Lease period factor—Controls how much ahead of the client the partner’s idea of the lease expiration can be. It is a multiple of the desired lease period used to update the partner when the main server informs it of a lease renewal. Possible values in the range of values are:
  - 1.0—Same as the lease period. The lease period factor should always be more than this value.
  - 1.5—The default and optimized factor. It is the lease period plus half again the lease, best used if the renewal period is 50% of the lease.
  - 2.0—Twice the lease period.

The lease period factor depends on the lease renewal period. If the renewal period is more than 50% of the lease, you must also increase the factor. The calculation is:

$$1 + \text{lease-renewal-percentage} = \text{lease-period-factor}$$

Thus, the usual renewal period of 50% might take the default (1 + 0.5 =) 1.5 lease period factor. A renewal period of 80% would more appropriately take a (1 + 0.8 =) 1.8 lease period factor.

You must define the lease period factor for the main DHCP server only. If defined for a partner, the main server ignores it, to enable duplicating the configurations through scripts.

Generally, if you enable failover on your DHCP server, you should not change the *failover-maximum-client-lead-time* or *lease-period-factor* attributes. However, you can do so explicitly:

- 
- Step 1** Reload the backup server to ensure that all data that the backup server has for the main server is up-to-date. Ideally, you should wait until both partners are stabilized, in NORMAL state, and any updates were exchanged. At least wait until the backup server completes its cache update process, as the log file indicates.
  - Step 2** Change the MCLT or lease period factor on the main server. The backup server ignores any MCLT that you configured on it, because it derives its MCLT value directly from the main. The default MCLT is 60 minutes and the default lease period factor is 1.5. Set the DHCP server attributes *failover-maximum-client-lead-time* and *lease-period-factor*.
  - Step 3** Stop the backup server.
  - Step 4** Reload the main server.
  - Step 5** Start the backup server.
- 

## Using the Failover Safe Period to Move Servers into PARTNER-DOWN State

One or both failover partners could potentially move into COMMUNICATIONS-INTERRUPTED state. Fortunately, they cannot issue duplicate addresses while in this state. However, having a server in this state over longer periods is not a good idea, because there are restrictions on what a server can do. The main server cannot re-allocate expired leases and the backup server can run out of addresses from its pool. COMMUNICATIONS-INTERRUPTED state was designed for servers to easily survive transient communication failures of a few minutes to a few days. A server might function effectively in this state for only a short time, depending on the client arrival and departure rate. After that, it would be better to move a server into PARTNER-DOWN state so it can completely take over the lease functions until the servers resynchronize.

There are two ways a server can move into PARTNER-DOWN state:

- User action—An administrator sets a server into PARTNER-DOWN state based on an accurate assessment of reality. The failover protocol handles this correctly.
- The failover safe period expires—When the servers run unattended for longer periods, they need an automatic way to enter PARTNER-DOWN state.

Network operators might not sense in time that a server is down or uncommunicative. Hence, the failover safe period, which provides network operators some time to react to a server moving into COMMUNICATIONS-INTERRUPTED state. During the safe period, the only requirement is that the operators determine that both servers are still running and, if so, fix the network communications failure or take one of the servers down before the safe period expires.

During this safe period, either server allows renewals from any existing client, but there is a major risk of possibly issuing duplicate addresses. This is because one server can suddenly enter PARTNER-DOWN state while the other is still operating. Because of this risk, the failover safe period is disabled by default. That is why it is best to enable the safe period only if, during a server failure, it is more important to get an address than risk receiving a duplicate one.

The length of the safe period is installation-specific, and depends on the number of unallocated addresses in the pool and the expected arrival rate of previously unknown clients requiring addresses. The safe period is typically 24 hours, although many environments can support periods of several days.

The number of extra addresses required for the safe period should be the same as the expected total of new clients a server encounters. This depends on the arrival rate of new clients, not the total outstanding leases. Even if you can only afford a short safe period, because of a dearth of addresses or a high arrival rate of new clients, you can benefit substantially by allowing DHCP to ride through minor problems that are fixable in an hour. There is minimum chance of duplicate address allocation, and re-integration after the solved failure is automatic and requires no operator intervention.

Here are some guidelines to follow to decide in using manual intervention or the safe period for transitioning to PARTNER-DOWN state:

- If your corporate policy is to have minimal manual intervention, set the safe period. Enable the DHCP attribute *failover-use-safe-period* to enable the safe period. Then, set the DHCP attribute *failover-safe-period* to set the duration (86400 seconds, or 24 hours, by default).
- If your corporate policy is to avoid conflict under any circumstances, then never let the backup server go into PARTNER-DOWN state unless by explicit command. Allocate sufficient addresses to the backup server so that it can handle new client arrivals during periods when there is no administrative coverage. You can set PARTNER-DOWN in two ways in Network Registrar:
  - On the View Failover Related Server page of the regional cluster Web UI, if the partner is in the Communications-interrupted failover state, you can click **Set Partner Down** in association with an input field for the PARTNER-DOWN date setting (see the “[Listing Related Servers for DHCP Servers](#)” section on page 5-3). This setting is initialized to the value of the *start-of-communications-interrupted* attribute. (In Normal Web UI mode, you cannot set this date to be an earlier value than the initialized date. In Expert Web UI mode, you can set this value to any date.) After clicking **Set Partner Down**, you return to the List Related Servers for DHCP Server page to view the result of the PARTNER-DOWN action.
  - Use the **dhcp setPartnerDown** command in the CLI, specifying the name of the partner server. This moves all the scopes running failover with the partner into PARTNER-DOWN state immediately, unless you specify a date and time with the command. This date and time should be when the partner was last known to be operational.

There are two conventions for specifying the date:

- *-num unit* (a time in the past), where *num* is a decimal number and *unit* is *s*, *m*, *h*, *d*, or *w* for seconds, minutes, hours, days or weeks respectively. For example, specify -3d for three days.

- Month (name or its first three letters), day, hour (24-hour convention), year (fully specified year or last two digits). This example notifies the backup server that its main server went down at 12 midnight on October 31, 2002:

```
nrcmd> server dhcp setPartnerDown dhcp2.example.com. -3d
nrcmd> server dhcp setPartnerDown dhcp2.example.com. Oct 31 00:00:00 2001
nrcmd> dhcp reload
```

**Note**

Wherever you specify a date and time in the CLI, enter the time that is local to the **nrcmd** process. If the server is running in a different time zone than this process, disregard the time zone where the server is running and use local time instead.

## Setting DHCP Request and Response Packet Buffers

The number of request buffers (set through the *max-dhcp-requests* DHCP attribute) sets the maximum number of simultaneous requests that the server can accept. The default value is 500, which is suitable for most deployments, but can be tuned to the capacity of the server. This capacity relates to the leasing rate and average latency of the leasing transaction. For example, if clients receive new leases every 250 milliseconds, then a request buffer value of 500 is sufficient for the server to respond to 2000 clients per second. (This assumes sufficient processing capacity to service clients at that rate.) A lower value would throttle the server's performance below this capacity. A higher value allows servicing a greater number of clients without retries during periods of burst load, but results in a higher average latency to each client. Average latency under two seconds is sufficient to properly service clients.

The number of response buffers (set through the *max-dhcp-responses* DHCP attribute) sets the maximum number of simultaneous requests that the server can complete by issuing a client response. When the network operates at a steady state, responses should track with the number of requests accepted. Because the same pool of response buffers serves for both lease and failover activity, when failover is enabled, the server adjusts the response buffer value to be at least four times the request buffers. This ensures that sufficient resources are available to process all pending client and failover activity simultaneously.

## Changing Polling Attributes

You can change some system defaults, such as the number of leases that the main server should send to the backup server, or the MCLT. See the [“Setting the Maximum Client Lead Time and Lease Period Factor” section on page 16-17](#). However, you need to change them on both servers.

On each server:

- Change the poll interval (DHCP attribute *failover-poll-interval*)—The interval that partners contact each other to confirm network connectivity. The default is 10 seconds.
- Change the poll timeout (DHCP attribute *failover-poll-timeout*)—Failover partners who cannot communicate for failover-poll-timeout seconds will conclude that they lost network connectivity, and change their operational states appropriately. The default is 60 seconds.

Generally, you should not have to change the failover-poll-timeout. It is intimately linked to the *failover-poll-interval* and is based on real world experience.

**Note**

To collect subnet utilization history for the failover pair, if you are configuring simple failover, disable individual polling of the main and backup DHCP servers, but enable failover pair polling by setting the failover pair attribute *poll-subnet-util-interval*, so as to collect one set of data from both servers.

## Setting the Network Discovery Attribute

If you enable failover on a UNIX system, you could set the *sms-network-discovery* attribute to enable the computing client os-type for leased addresses—This can help if you have a Windows partner server and want to use the **dhcp updateSms** CLI command on it.

## Changing Failover Server Roles

**Caution**

Be careful when you change the role of a failover server. Remember that all address states in a scope are lost from a server if it is ever reloaded without that scope in its configuration.

## Making Nonfailover Servers Failover Mains

You can update an existing installation and increase the availability of the DHCP service it offers. You can use this procedure only if the original server never participated in failover.

- 
- Step 1** Install Network Registrar on the original server and ensure that it operates correctly after the installation.
  - Step 2** Install Network Registrar on the machine that is to be the backup server. Note the machine's DNS name.
  - Step 3** Enable failover on the original server. Use the DNS name of the recently installed backup server. See the [“Simple Failover” section on page 16-1](#).
  - Step 4** Reload the main server. It should go into PARTNER-DOWN state and stay there. It cannot locate the backup server, because it is not yet configured. There should be no change in main server operation at this point.
  - Step 5** Duplicate the main server's configuration on the backup server, including scopes (including secondary), policies, and client-classes. If you use client-classes, make sure the clients are entered into each cluster or that each server can access an LDAP database with the client data.
  - Step 6** Enable failover on the backup server. Be sure to define the main server.
  - Step 7** Reconfigure all the operational BOOTP relays to forward broadcast DHCP packets to both the main and backup server.
  - Step 8** Reload the backup server.
- 

After you complete these steps:

1. The backup server detects the main server and moves into RECOVER state.
2. The backup server refreshes its stable storage with the main server's lease data and, when complete, moves into RECOVER-DONE state.
3. The main server moves into NORMAL state.
4. The backup server moves into NORMAL state.
5. The backup server uses a pool request to ask the main server for addresses to allocate if communication is interrupted.
6. After allocating these addresses, the main server sends this data to the backup server.

## Replacing Servers Having Defective Storage

If a failover server loses its stable storage (hard disk), you can replace the server and have it recover its state information from its partner:

- 
- Step 1** Determine which server lost its stable storage.
  - Step 2** Use the **Set Partner Down** feature of the Web UI or the **dhcp setPartnerDown** command in the CLI to tell the other server that its partner is down. If you do not specify a time, Network Registrar uses the current time.
  - Step 3** When the server is again operational, re-install Network Registrar.
  - Step 4** Duplicate the configuration on the server from its partner.
  - Step 5** Reload the replacement server.
- 

After you complete these steps:

1. The recovered server moves into RECOVER state.
2. Its partner sends it all its data.
3. The server moves into RECOVER-DONE state when it reaches its maximum client lead time (and any time set for *failover-recover*).
4. Its partner moves into NORMAL state.
5. The recovered server moves into NORMAL state. It can request addresses, but can allocate few new ones, because its partner already sent it all its previously allocated addresses.

## Removing Backup Servers and Halting Failover Operation

There are times when you might need to remove the backup server and halt all failover operations.

- 
- Step 1** On the backup server, remove all the scopes that were designated as a backup to the main server.
  - Step 2** On the main server, remove the failover capability from those scopes that were main for the backup server, or disable failover server-wide if that is how it was configured.
  - Step 3** Reload both servers.
- 

## Adding Main Servers to Existing Backup Servers

You can use an existing backup server for a main server.

- 
- Step 1** Duplicate the main server's scopes, policies, and other configurations on the backup server.
  - Step 2** Configure the main server to enable failover and point to the backup server.
  - Step 3** Configure the backup server to enable failover for the new scopes that point to the new main server.

- Step 4** Reload both servers. Network Registrar performs the same steps as those described in the [“Making Nonfailover Servers Failover Mains”](#) section on page 16-21.
- 

## Configuring Failover on Multiple Interface Hosts

If you plan to use failover on a server host with multiple interfaces, you must explicitly configure the local server's name or address. This requires an additional command. For example, if you have a host with two interfaces, serverA and serverB, and you want to make serverA the a main failover server, you must define serverA as the failover-main-server before you set the backup server name (external serverB). If you do not do this, failover might not initialize correctly and tries to use the wrong interface.

Set the DHCP server properties *failover-main-server* and *failover-backup-server*.

With multiple interfaces on one host, you must specify a host name that points to only one address or A record. You cannot set up your servers for round-robin support.

## Supporting BOOTP Clients in Failover

You can configure scopes to support two types of BOOTP clients—static and dynamic.

### Static BOOTP

You can support static BOOTP clients using DHCP reservations. When you enable failover, remember to configure both the main and the backup server with identical reservations.

### Dynamic BOOTP

You can enable dynamic BOOTP clients by enabling the *dynamic-bootp* attribute on a scope. When using failover, however, there are additional restrictions on address usage in such scopes, because BOOTP clients get permanent addresses and leases that never expire.

When a server whose scope does not have the *dynamic-bootp* option enabled goes to PARTNER-DOWN state, it can allocate any available (unassigned) address from that scope, whether or not it was initially available to any partner. However, when the *dynamic-bootp* option is set, each partner can only allocate its own addresses. Consequently, scopes that enable the *dynamic-bootp* option require more addresses to support failover.

When using dynamic BOOTP:

- Segregate dynamic BOOTP clients to a single scope. Disable DHCP clients from using that scope by disabling the *dhcp* attribute on the scope.
- Set the *failover-dynamic-bootp-backup-percentage* DHCP server attribute to allocate a greater percentage of addresses to the backup server for this scope, as much as 50 percent higher than a regular backup percentage.

## Configuring BOOTP Relays

The Network Registrar failover protocol works with BOOTP relay (also called IP helper), a router capability that supports DHCP clients that are not locally connected to a server.

If you use BOOTP relay, ensure that the implementations point to both the main and backup servers. If they do not and the main fails, clients are not serviced, because the backup cannot see the required packets. If you cannot configure BOOTP relay to forward broadcast packets to two different servers, configure the router to forward the packets to a subnet-local broadcast address for a LAN segment, which could contain both the main and backup servers. Then, ensure that both the main and backup servers are on the same LAN segment.

## DHCPLEASEQUERY and Failover

To accommodate DHCPLEASEQUERY messages sent to a DHCP failover backup server when the master server is down, the master server must communicate the *relay-agent-info* (82) option values to its partner server. To accomplish this, the master server uses DHCP failover update messages.


## Troubleshooting Failover

This section describes how to avoid failover configuration mistakes, monitor failover operations, and detect and handle network problems.

## Monitoring Failover Operations

You can examine the DHCP server log files on both partner servers to verify your failover configuration.

You can make a few important log and debug settings to troubleshoot failover. Set the DHCP log settings to *failover-detail* the number and detail of failover messages logged. To ensure that previous messages do not get overwritten, add the *failover-detail* attribute to the end of the list. Use the *no-failover-conflict* attribute to inhibit logging server failover conflicts, or the *no-failover-activity* attribute to inhibit logging normal server failover activity. Then, reload the server.

You can also isolate misconfigurations more easily by clicking the Related Servers icon () on the Manage DHCP Server or List DHCP Failover Pairs page in the Web UI (see the “[Listing Related Servers for DHCP Servers](#)” section on page 5-3), or use the `dhcp getRelatedServers` command in the CLI.

## Detecting and Handling Network Failures

Table 16-4 describes some symptoms, causes, and solutions for failover problems.

**Table 16-4** *Detecting and Handling Failures*

| Symptom                          | Cause                                                                         | Solution                                           |
|----------------------------------|-------------------------------------------------------------------------------|----------------------------------------------------|
| New clients cannot get addresses | A backup server is in COMMUNICATIONS-INTERRUPTED state with too few addresses | Increase the backup percentage on the main server. |

**Table 16-4** *Detecting and Handling Failures (continued)*

| Symptom                                                                                                                                                          | Cause                                                                                                                                               | Solution                                                                                                                                                                                                                                            |
|------------------------------------------------------------------------------------------------------------------------------------------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Error messages about mismatched scopes                                                                                                                           | There are mismatched scope configurations between partners                                                                                          | Reconfigure your servers.                                                                                                                                                                                                                           |
| Log messages about failure to communicate with partner                                                                                                           | Server cannot communicate with its partner                                                                                                          | Check the status of the server.                                                                                                                                                                                                                     |
| Main server fails. Some clients cannot renew or rebind leases. The leases expire even when the backup server is up and possibly processing some client requests. | Some BOOTP relay (ip-helper) was not configured to point at both servers; see the <a href="#">“Configuring BOOTP Relays”</a> section on page 16-24. | <ul style="list-style-type: none"> <li>Reconfigure BOOTP relays to point at both main and backup server</li> <li>Run a fire drill test—Take the main server down for a day or so and see if your user community can get and renew leases</li> </ul> |
| SNMP trap: other server not responding                                                                                                                           | Server cannot communicate with its partner                                                                                                          | Check the status of the server.                                                                                                                                                                                                                     |
| SNMP trap: dhcp failover configuration mismatch                                                                                                                  | Mismatched scope configurations between partners                                                                                                    | Reconfigure your servers.                                                                                                                                                                                                                           |
| Users complain that they cannot use services or system as expected                                                                                               | Mismatched policies and client-classes between partners                                                                                             | Reconfigure partners to have identical policies; possibly use LDAP for client registration if currently registering clients directly in partners.                                                                                                   |

